

REPORT DOCUMENTATION PAGE				<i>Form Approved</i> OMB No. 0704-0188	
The public reporting burden for this collection of information is estimated to average 1 hour per response, including the time for reviewing instructions, searching existing data sources, gathering and maintaining the data needed, and completing and reviewing the collection of information. Send comments regarding this burden estimate or any other aspect of this collection of information, including suggestions for reducing this burden, to Department of Defense, Washington Headquarters Services, Directorate for Information Operations and Reports (0704-0188), 1215 Jefferson Davis Highway, Suite 1204, Arlington, VA 22202-4302. Respondents should be aware that notwithstanding any other provision of law, no person shall be subject to any penalty for failing to comply with a collection of information if it does not display a currently valid OMB control number. PLEASE DO NOT RETURN YOUR FORM TO THE ABOVE ADDRESS.					
1. REPORT DATE (DD-MM-YYYY) APRIL 2016		2. REPORT TYPE CONFERENCE PAPER (Post Print)		3. DATES COVERED (From - To) JUN 2014 – SEP 2015	
4. TITLE AND SUBTITLE EVOLVING AND CONTROLLING PERIMETER, RENDEZVOUS, AND FORAGING BEHAVIORS IN A COMPUTATION-FREE ROBOT SWARM				5a. CONTRACT NUMBER IN-HOUSE	
				5b. GRANT NUMBER NA	
				5c. PROGRAM ELEMENT NUMBER 62788F	
6. AUTHOR(S) Daniel Brown and Matthew Johnson				5d. PROJECT NUMBER S2MA	
				5e. TASK NUMBER MI	
				5f. WORK UNIT NUMBER IH	
7. PERFORMING ORGANIZATION NAME(S) AND ADDRESS(ES) Air Force Research Laboratory/Information Directorate Rome Research Site/RISC 525 Brooks Road Rome NY 13441-4505				8. PERFORMING ORGANIZATION REPORT NUMBER N/A	
9. SPONSORING/MONITORING AGENCY NAME(S) AND ADDRESS(ES) Air Force Research Laboratory/Information Directorate Rome Research Site/RISC 525 Brooks Road Rome NY 13441-4505				10. SPONSOR/MONITOR'S ACRONYM(S) AFRL/RI	
				11. SPONSORING/MONITORING AGENCY REPORT NUMBER AFRL-RI-RS-TP-2016-005	
12. DISTRIBUTION AVAILABILITY STATEMENT APPROVED FOR PUBLIC RELEASE; DISTRIBUTION UNLIMITED. PA Case Number: 88ABW-2015-5024 DATE CLEARED: 16 OCT 2015					
13. SUPPLEMENTARY NOTES Proceedings 9 th European Alliance for Innovation (EAI) International Conference on Bio-Inspired Information and Communication Technologies, New York City, NY, 03-05 Dec 2015. One or more of the authors is a U.S. Government employee working within the scope of their Government job; This is a work of the United States Government and is not subject to copyright protection in the United States.					
14. ABSTRACT Designing and controlling simple collective robot behaviors often requires complex range and bearing sensors and peer-to-peer communication strategies. Recent work studying swarms robots that have no computational power has shown that complex behaviors such as aggregation and clustering items can be produced from extremely simple control policies and sensing capability. We extend previous work on computation-free swarm behaviors and show that it is possible to evolve simple control policies to form a perimeter around a target, rendezvous to a specific location, and perform foraging. We also demonstrate that simple manipulations of the environment provide a form of stigmergic control, whereby these collective behavior can be controlled. The robustness and expressiveness of these behaviors, combined with the simple requirements for control and sensing, demonstrate the feasibility of implementing swarm behaviors at small scales or in extreme environments.					
15. SUBJECT TERMS Stochastic physical search, Planning under uncertainty, Multiagent optimization, k-Agent sufficiency					
16. SECURITY CLASSIFICATION OF:			17. LIMITATION OF ABSTRACT UU	18. NUMBER OF PAGES 5	19a. NAME OF RESPONSIBLE PERSON STEVEN LOSCALZO
a. REPORT U	b. ABSTRACT U	c. THIS PAGE U			19b. TELEPHONE NUMBER (Include area code) N/A

Evolving and Controlling Perimeter, Rendezvous, and Foraging Behaviors in a Computation-Free Robot Swarm

Matthew Johnson
Air Force Research Laboratory
26 Electronic Parkway
Rome, NY 13441
matthew.johnson.151@us.af.mil

Daniel S. Brown
Air Force Research Laboratory
26 Electronic Parkway
Rome, NY 13441
daniel.brown.81@us.af.mil

ABSTRACT

Designing and controlling simple collective robot behaviors often requires complex range and bearing sensors and peer-to-peer communication strategies. Recent work studying swarms robots that have no computational power has shown that complex behaviors such as aggregation and clustering items can be produced from extremely simple control policies and sensing capability. We extend previous work on computation-free swarm behaviors and show that it is possible to evolve simple control policies to form a perimeter around a target, rendezvous to a specific location, and perform foraging. We also demonstrate that simple manipulations of the environment provide a form of stigmergic control, whereby these collective behavior can be controlled. The robustness and expressiveness of these behaviors, combined with the simple requirements for control and sensing, demonstrate the feasibility of implementing swarm behaviors at small scales or in extreme environments.

Categories and Subject Descriptors

I.2.11 [Artificial Intelligence]: Distributed Artificial Intelligence—*multiagent systems, coherence and coordination*;
I.2.9 [Artificial Intelligence]: Robotics—*intelligent vehicles*

Keywords

swarm robotics, evolutionary algorithms, computation-free robot, controlling collective behaviors

1. INTRODUCTION

Flocks of birds, schools of fish, and colonies of ants, bees, termites exhibit a remarkable robustness and resilience, despite the limited capabilities of each individual. Recently research into bio-inspired swarm robotics has been gaining popularity due to the low-cost, robust, redundant, and distributed nature of swarms [3]. Potential applications for robot swarms include, search and rescue, construction, and

chemical spill clean-up, as well as nano-medical applications such as finding tumors [10]. Many of these applications would benefit from simple, cheap, disposable swarms of robots that can accomplish these tasks quickly and without much human supervision.

While there has been a lot of work on different swarming algorithms and technologies, many still require localization, mapping, complex coordination algorithms, and precise identification of neighboring robots' orientations and relative positions. This often results in swarm behaviors that are interesting but extremely difficult to implement on actual robotic platforms. For swarm applications in the nano-medical field, developing collective behaviors that use extremely simple controllers and sensors is especially important if these behaviors have an hope of being implemented on nano-robots[11].

Recently, Gauci et al. have shown that swarms of robots so dumb that they have no computational power—they can't even add or subtract, and have no memory can still collectively solve canonical multi-robot problems such as aggregation [6], and simple object clustering[5]. There are several key benefits to researching the capabilities of extremely dumb robot swarms: (1)the dumber the robot, the cheaper and more disposable it is, (2) the simpler the control algorithm the easier to implement on real robots[8, 6], and (3) even teams of smart robots may need a "Plan B" consisting of simple robust algorithms that require only the most basic capabilities in case of malfunctions and failure.

We extend the work of Gauci et al. by showing that many other interesting behaviors can be achieved using swarms of computation-free robots. Our work starts with the simple robot model proposed in [6] and adds a form of stigmergic control by changing the environment to expand the possible behaviors and control collective behaviors. We investigate what behaviors are possible given limited control over the placement of a small number of objects in the swarm's environment. We use a genetic algorithm approach to design these swarm behaviors by first defining a fitness function that describes a desired collective behavior, and then searching the space of simple controllers that best achieve this behavior. This approach allows us to evolve successively better controllers using a robot simulator to evaluate potential controllers. We present successful results on three behaviors that are possible using very simple sensors and controllers: forming a perimeter, rendezvous, and foraging. We additionally show that simple manipulations of the environment allow these behaviors to be controlled.

Permission to make digital or hard copies of all or part of this work for personal or classroom use is granted without fee provided that copies are not made or distributed for profit or commercial advantage and that copies bear this notice and the full citation on the first page. To copy otherwise, to republish, to post on servers or to redistribute to lists, requires prior specific permission and/or a fee.

2. RELATED WORK

Trianni et al. evolved a neural network based controller that performs aggregation using swarms of S-bots [13]. However, each S-bot uses eight infrared proximity sensors, three microphones, three sensors for detecting connections on the body and a gripper sensor. Baldassarre et al. evolved a controller that aggregates a group of robots and then moves them towards a light source [1]. Their controller utilizes a neural network that takes in eight infrared proximity sensors readings, four directional light source sensors readings, and four directional sound sensors readings as control inputs. Gauci et al. introduced the concept of robots that can't compute [4]. They showed that a simple reactive controller could be used to allow a swarm of computation-free robots with a single line-of-sight sensor to perform aggregation [6] and clustering [5].

Other work has looked at controlling collective behaviors. Rubenstein et al. [12] studied how to collectively transport items using a simple control signals and behaviors. Others have looked at controlling more complex collective transport problems [2] or using termite-inspired stigmergic control to build complex structures [14]. However, none of this work considers the extreme conditions of a single line-of-sight sensor and zero computation.

3. PROBLEM FORMULATION

This paper investigates what collective behaviors are possible given a swarm of extremely simple robots operating in a simple environment. For our experiments we consider a circularly bounded 2D environment that is homogeneous and contains no obstacles. Throughout this space n circular robots are randomly distributed and randomly rotated such that each robot faces a random direction. Robots learn to interact with immovable targets, movable objects, and other robots to achieve global behaviors. All entities (robots, targets, and objects) are rigid such that no two entities can occupy the space at the same time.

We define simple robots as agents that are memoryless, cannot perform computations and have limited input/output capabilities. Specifically we look at robots that are only equipped with a line of sight sensor and two wheels for differential drive. The line of sight sensor can only detect the presence or absence of objects and outputs a trinary value where $s = 2$ corresponds to a target or object in line of sight, $s = 1$ corresponds to a robot in line of sight, and $s = 0$ corresponds to nothing in line of sight.

Simple robots are reactive in nature because they cannot remember past input or actions. As a result simple robot controllers can be to a sequential series of if-statements that assign left and right wheel velocities based on current sensor readings. This controller can be represented as a set of six wheel velocities

$$V = [v_{l0}, v_{r0}, v_{l1}, v_{r1}, v_{l2}, v_{r2}]$$

where v_{l0}/v_{r0} are the velocities of the left/right wheel when there is no robot in the sensors current line of sight, v_{l1}/v_{r1} are the respective velocities when a robot is within the line of sight of the sensor and v_{l2}/v_{r2} are the respective velocities when a target or object is within the line of sight of the sensor. Velocities are normalized such that $v = [-1, 1]$ where 1 corresponds to a wheel spinning forward at full speed and -1 corresponds to a wheel spinning backwards at full speed.




Figure 1: Robot representation and corresponding controller.

4. BEHAVIORS

This section explores several global behaviors learned using evolutionary optimization techniques. We discover global behaviors by optimizing a universal robotic controller according to a behavior dependent fitness function. Each potential robot controller is evaluated by running a swarm simulation and calculating the fitness function at every time step to generate a fitness score, which is given by

$$U(V) = \sum_{t=0}^{T-1} tu(t)$$

where T is the number of time steps in the simulation and $u(t)$ is the fitness function. Multiplying the fitness function by the time step rewards controllers that achieve desired behavior quickly. The robot controllers are optimized using the average fitness score over multiple simulations to reduce the effect of noise.

All simulations are run on the Enki 2.0 robot simulator, which is able to simulate hundreds of robots in a 2D environment in faster than real time [9]. For our experiments the simulation physics are updated 100 times per second and the robot controller is updated 10 times per second. Robots are simulated using Enki's Epuck model which have a diameter of 7.4 cm, inter-wheel distance of 5.1 cm, and weight of 152 g. Targets and objects are simulated as cylinders with a diameter of 10 cm using Enki's physical object model. Objects have a mass of 35g and a coefficient of friction of 0.58. Targets have a sufficiently large mass and coefficient of friction to ensure that they are immobile.

Robot controllers are optimized using the Covariance Matrix Adaption Evolution Strategy (CMA-ES) [7]. This genetic optimization technique uses the variance of each gene to generate mutations between generations. Earlier work by Gauci et al. has shown that CMA-ES can effectively optimize simple robotic controllers [5]. CMA-ES optimizes across all real numbers, which can result in genes out of normalized range. To avoid this we constrain genes by applying the following sigmoid function

$$v = \frac{1 - e^{-x}}{1 + e^{-x}}$$

where x is a gene optimized by CMA-ES. For our experiments we utilized the following CMA-ES parameters: population size of 13, initial step size, $\sigma(0) = 0.72$, and starting controller of $V = [0, 0, 0, 0, 0, 0]$.

4.1 Aggregating to a Target

We first investigate what is possible when a single stationary target is placed in the environment. In this behavior




Figure 2: Perimeter formation around a dynamic target.

robots are initially distributed randomly throughout an environment and over time form a group around a randomly positioned target. Our fitness function rewards global behaviors that minimize the total distance between each robot and the target. Let $p_i(t)$ represent the position of robot i at time step t and p_{target} represent the position of the target. Then the rendezvous fitness function is given by

$$u_{rendezvous}(t) = - \sum_{i=0}^{n-1} \|p_i(t) - p_{target}\|^2$$

where $\|\cdot\|$ is the Euclidean norm. This fitness function rewards solutions where the robots are close to the objects locations.

4.1.1 Perimeter Formation

The resulting controller that is evolved forms a perimeter around the target. The evolved controller is

$$V = [1.0, 0.37, 1.0, 1.0, -1.0, 0.83].$$

This solution results in the robots aggregating to the target and forming a perimeter around the target. We also experimented with changing the location of the target mid-simulation. The results are shown in Figure 2. The robots converge to the target and form a perimeter. When the target is placed in a new location, the entire swarm quickly moves to the new location and reforms the perimeter. This behavior is very robust and is automatic—the robots are purely reactive so they can be controlled simply by changing the environment, removing the need to broadcast information to the swarm or have additional control logic.

4.1.2 Rendezvous

We are also interested in having every member of the swarm gather as close as possible to the target, rather than just circle around it. Rendezvous is an important behavior for swarms because it sets the stage for more complicated behaviors by assembling a group of robots to a specific desired location. We first tried to find a controller for the rendezvous problem using the fitness function described above; however, all trials resulted in controllers in which robots would form a circle around a target.

To solve this problem we seeded the starting controller with an aggregation solution from Gauci’s et al. earlier work

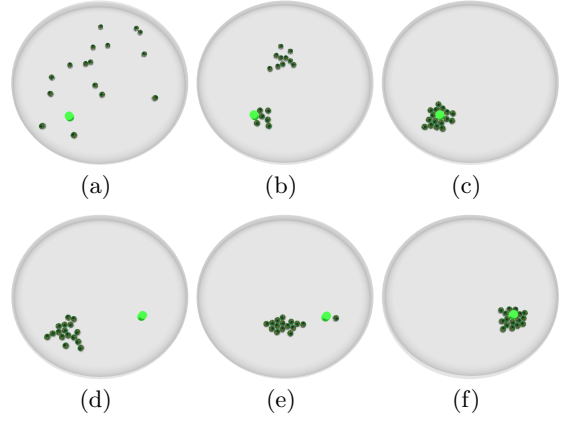


Figure 3: Rendezvous to dynamic target.

[6]. Using the seeded optimization strategy we evolved the following controller

$$V = [-0.72, -1.00, 1.00, -1.00, 0.99, 1.00]$$

which results in the rendezvous behavior. Using this controller robots closest to the target move towards it with a slightly elliptical path, while other robots form clusters elsewhere. Over time the cluster of robots around the target pulls the extra clusters toward the target cluster. When the clusters are close enough they merge into a single cluster around the target. Snapshots of the resulting behavior are shown in Figure 3. Similar to the perimeter formation, this behavior can also be controlled by simply moving the target.

4.2 Foraging

In this behavior objects and robots are distributed randomly throughout the environment and the robots must gather the objects to a specified target location. Earlier work by Gauci et al. found an optimal controller for clustering objects, we extend their work by showing that this controller can be used for foraging [5]. The clustering fitness function rewards global behaviors that minimize the total distance between each object and center of the cluster of objects. Let $o_i(t)$ represent the position of object i at time step t and $\bar{o}(t)$ represent the center of the object cluster. Then the fitness function is given by

$$u_{clustering}(t) = - \sum_{i=0}^{m-1} \|o_i(t) - \bar{o}\|^2$$

where m is the number of objects. Using the clustering fitness function we evolved the following controller

$$V = [0.72, 1.00, 0.40, 0.31, 0.53, -1.00]$$

which causes the robots to circle around the objects and slowly nudge them into a central point as they pass.

The foraging behavior occurs when we place one or several fixed targets in the environment. Figure 4 shows the classic foraging problem where there is a “nest” location (shown in green) where all of the items must be gathered. Figure 5 shows an alternative foraging scheme where multiple stationary targets are placed in the environment. The convex hull of these targets defines the region into which the objects will be harvested. Similar to the previous behaviors, the foraging behavior can be controlled simply by changing




Figure 4: Foraging




Figure 5: Dynamic foraging into specified convex hull.

the location of the targets. The robots will then move the items to the new desired location, as shown in Figure 5.

5. CONCLUSIONS AND FUTURE WORK

A large amount of research has been dedicated to developing multi-agent systems that perform complex behaviors. We show that swarms of robots that can't compute can perform complex behaviors such as rendezvous to a desired location, simple perimeter monitoring of a desired location, and foraging in changing environments. Our results demonstrate that complex behaviors can be evolved from simple interactions between agents and that these behaviors can be controlled during execution by simply changing the environment. We also note that these behaviors are so simple that they could simply be hardwired, requiring no computational capabilities. We believe that this research is an important step towards swarm behaviors that can be easily implemented in hardware and produced at small, maybe even nano-scale. In the future we plan to apply these behaviors to actual robots, explore virtual targets and other

forms of stigmergic control, and more rigorously explore the space of possible behaviors given our computation-free assumptions.

6. REFERENCES

- [1] G. Baldassarre, S. Nolfi, and D. Parisi. Evolving mobile robots able to display collective behaviours. *Artificial Life*, 9(3):255–267, 2003.
- [2] A. Becker, C. Ertel, and J. McLurkin. Crowdsourcing swarm manipulation experiments: A massive online user study with large swarms of simple robots. In *Robotics and Automation (ICRA), 2014 IEEE International Conference on*, pages 2825–2830. IEEE, 2014.
- [3] M. Brambilla, E. Ferrante, M. Birattari, and M. Dorigo. Swarm robotics: a review from the swarm engineering perspective. *Swarm Intelligence*, 7(1):1–41, 2013.
- [4] M. Gauci, J. Chen, T. J. Dodd, and R. Groß. Evolving aggregation behaviors in multi-robot systems with binary sensors. In *Distributed Autonomous Robotic Systems*, pages 355–367. Springer, 2014.
- [5] M. Gauci, J. Chen, W. Li, T. J. Dodd, and R. Groß. Clustering objects with robots that do not compute. In *Proceedings of the 2014 international conference on Autonomous agents and multi-agent systems*, pages 421–428. International Foundation for Autonomous Agents and Multiagent Systems, 2014.
- [6] M. Gauci, J. Chen, W. Li, T. J. Dodd, and R. Groß. Self-organized aggregation without computation. *The International Journal of Robotics Research*, page 0278364914525244, 2014.
- [7] N. Hansen and A. Ostermeier. Completely derandomized self-adaptation in evolution strategies. *Evolutionary computation*, 9(2):159–195, 2001.
- [8] N. Jakobi. Half-baked, ad-hoc and noisy: Minimal simulations for evolutionary robotics. In *Fourth European Conference on Artificial Life*, pages 348–357. MIT press, 1997.
- [9] S. Magnenat, M. Waibel, and A. Beyeler. Enki: The fast 2d robot simulator, 2007.
- [10] S. M. Moghimi, A. C. Hunter, and J. C. Murray. Nanomedicine: current status and future prospects. *The FASEB Journal*, 19(3):311–330, 2005.
- [11] A. A. Requicha. Nanorobots, nems, and nanoassembly. *Proceedings of the IEEE*, 91(11):1922–1933, 2003.
- [12] M. Rubenstein, A. Cabrera, J. Werfel, G. Habibi, J. McLurkin, and R. Nagpal. Collective transport of complex objects by simple robots: theory and experiments. In *Proceedings of the 2013 international conference on Autonomous agents and multi-agent systems*, pages 47–54. International Foundation for Autonomous Agents and Multiagent Systems, 2013.
- [13] V. Trianni, R. Grob, T. Labella, E. Sahin, and M. Dorigo. Evolving aggregation behaviors in a swarm of robots. In *Advances in artificial life*, pages 865–874. Springer, 2003.
- [14] J. Werfel, K. Petersen, and R. Nagpal. Designing collective behavior in a termite-inspired robot construction team. *Science*, 343(6172):754–758, 2014.